



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Inferring low-dimensional microstructure representations using convolutional neural networks

Nicholas Lubbers, Turab Lookman, and Kipton Barros

Phys. Rev. E **96**, 052111 — Published 9 November 2017

DOI: [10.1103/PhysRevE.96.052111](https://doi.org/10.1103/PhysRevE.96.052111)

Inferring low-dimensional microstructure representations using convolutional neural networks

Nicholas Lubbers,^{1,2,*} Turab Lookman,² and Kipton Barros²

¹*Department of Physics, Boston University, Boston, MA*

²*Theoretical Division and CNLS, Los Alamos National Laboratory, Los Alamos, NM 87545*

(Dated: September 29, 2017)

We apply recent advances in machine learning and computer vision to a central problem in materials informatics: The statistical representation of microstructural images. We use activations in a pre-trained convolutional neural network to provide a high-dimensional characterization of a set of synthetic microstructural images. Next, we use manifold learning to obtain a low-dimensional embedding of this statistical characterization. We show that the low-dimensional embedding extracts the parameters used to generate the images. According to a variety of metrics, the convolutional neural network method yields dramatically better embeddings than the analogous method derived from two-point correlations alone.

I. INTRODUCTION

A central problem in materials design is the analysis, characterization, and control of materials microstructure. Microstructure is generated by non-equilibrium processes during the formation of the material and plays a large role in the bulk material’s properties [1–5]. In recent years, machine learning and informatics based approaches to materials design have generated much interest [6–10]. Effective statistical representation of microstructure has emerged as an outstanding challenge [11–13].

Standard approaches begin with an n -point expansion, and typically truncate at the pair correlation level [14–17]. Pair correlations can capture information such as the scale of domains in a system, but miss higher order complexities such as the detailed shape of domains or relative orientation of nearby domains [18–21]. Three-point correlations (and successors) quickly become computationally infeasible, as the number of n -point correlations scales exponentially with n . Furthermore, they are not tailored to capture the statistical information of interest. Much current work involves deploying a set of modified two-point correlations to better capture certain microstructural features [16, 22–25].

Independently, researchers in machine learning for computer vision have been developing powerful techniques to analyze image content [26–30]. Deep Convolutional Neural Networks (CNNs) have emerged as a particularly powerful tool for image analysis [31]. Of particular interest to materials microstructure analysis is literature regarding *texture* reconstruction and modeling [32–36]; in this context a texture is an image with roughly translation invariant statistics. Indeed, Gatys et al. have recently demonstrated that correlations between CNN activations capture the statistics of textures exceptionally well [37, 38].

Here, we apply the Gatys et al. CNN *texture vector* representation of image statistics to the problem of char-

acterizing materials micrographs. The texture vector is a statistical representation of an image derived from the processing of a pre-trained CNN; we use the Visual Geometry Group’s VGG-19 network [28], which has been trained to classify 1.2 million natural images [39]. We demonstrate that the texture vectors generated using the VGG-19 network can capture complex statistical correlations visible in microstructure images in Fig. 1. Specifically, we use the CNN texture vector to characterize an original microstructural image, and then generate a new, random image constrained to the same statistics. It is remarkable that using only a single original image, the algorithm generates texture images nearly indistinguishable to the eye. In the case of materials micrographs, where data can be expensive to collect, the ability for a method perform well on small datasets is crucial. Our approach can be considered one of *transfer learning*, *i.e.*, the application of a model trained on one problem to achieve results on a different, but related problem. The fidelity in Fig. 1 motivates us to pursue the CNN texture vector as a tool for mapping relationships between materials processing, microstructure, and properties.

As a step towards microstructure analysis, we demonstrate and quantify the ability of the CNN texture vector to extract hidden information from a collection of synthetic texture images. Images in our datasets are generated under different “processing” conditions, *i.e.* a few generating parameters. The goal is to extract a compact statistical representation of the dataset that captures the relevant statistics associated with the hidden generating parameters—that is, to back out the processing parameters of micrographs directly from the images alone. Our use of synthetic data, with a tightly controlled ground truth, provides us quantitative measures of error in the unsupervised learning process.

Despite being very high-dimensional, the CNN texture vector offers a good notion of abstract distance between texture images: The closer the generating parameters, the smaller the distance between texture vectors. We use *manifold learning* to embed each image as a point in a low-dimensional space such that, ideally, the embed-

* nlubbers@bu.edu




FIG. 1: (Color online) Texture synthesis of materials microstructures using the CNN algorithm from [37]. The CNN synthesizes each “Reconstruction” image from a single “Original” image. (Image attributions listed in Appendix A.)

ded distances match the texture vector distances. The structure of the embedded points then reveals information about the set of images. For our synthetic dataset, we show a simple relationship between the embedded coordinates and the generating parameters. More broadly, dimensionality reduction techniques may serve as the basis for characterization of materials properties that are controlled by complex materials microstructures. A recent example is the work of Ref. 40, which uses a method

similar to ours to map the space of ultrahigh carbon steel microstructures.

Our approach applies *unsupervised* learning, a pattern-discovery framework to seek new aspects of microstructure without using labeled data. This approach is applicable to problems where the ground truth is unknown, *e.g.*, the forensic analysis of microstructures. Several recent applications of machine learning to microstructure have used *supervised* learning algorithms [41–47] such as

support vector machines and classification trees, which make inferences based on labeled data. Like our work, Ref. 46 uses image features extracted from CNNs to aid microstructure analysis.

The remainder of the paper is organized as follows: Section II gives background on recent CNN architectures for image recognition, and specifics of the VGG network. Section III details our algorithms for statistical microstructure analysis and Sec. IV evaluates their accuracy on test datasets. Section V provides discussion and interpretation of our results, and we conclude in Sec. VI.

II. REVIEW OF CONVOLUTIONAL NEURAL NETWORKS

CNNs have emerged in recent years as state-of-the-art systems for computer vision tasks [26–29]. They form a modern basis for image recognition and object detection tasks, and in some cases now outperform humans [30, 48].

The basic computational structure is that of a many-layered (*i.e.*, *deep*) artificial neural network. For a brief overview, see Ref. 31; for a comprehensive text, see Ref. 49. There are a great variety of deep neural network architectures; here we first focus on the core components. Each layer in the network contains many computational units (*neurons*). Each neuron takes a set of inputs from the previous layer and computes a single output (the *activation*) to be used as an input in the next layer. Each neuron’s activation is constructed as follows: First, the set of inputs is linearly combined into a scalar using a set of *weights* and shifted using a *bias*. To this sum, the neuron applies a simple nonlinear map, the *activation function*, to generate its activation.

In the learning phase, the network is *trained* by iteratively tuning the weights and biases so that the network better performs a task. Performance of the network is quantified by a scalar *objective function*. Commonly, a network is trained by *supervised learning*, in which the network learns a mapping from inputs to outputs using a database of training examples. In this case, the objective function is a measure of error in the network output summed over all examples of the training set. The objective function is often differentiable and optimized via *stochastic gradient descent*.

A CNN is a specific type of artificial neural network which is useful for processing data on a spatial and/or temporal grid. The *convolutional layers* in CNNs impose strong restrictions on the structure of weights: Each layer consists of a bank of trainable filters (sometimes called *kernels*) that are convolved with activations from the previous layer. The convolution outputs are called *activation maps*. This technique of constraining and reusing weights is called *weight tying*. Note that the convolutional structure preserves spatial locality: The activation maps at each convolutional layer are interpretable as images. As in a plain artificial neural network, each pixel in the output image is passed through a nonlinear

activation function. CNNs also commonly include *pooling layers* that effectively coarse-grain the image plane. These layers operate by taking a statistic over a small region of the image plane, such as the maximum of a feature’s activations in a 2×2 pixel region. Importantly, the convolutional and pooling layers process the input image in a (nearly) translation equivariant way. This directly encoded translational symmetry is designed to match that of natural images, which as a distribution exhibit repeated patterns centered at a variety of locations. By alternating between sets of convolution and pooling layers, CNNs are able to develop sensitivity to very complex correlations over large length scales, which underlies their strong performance on image recognition tasks.

As in Gatys *et al.* [37, 38], our work begins with a normalized version of the Visual Geometry Group’s VGG-19 network [28] already trained to classify natural images. The VGG-19 network placed first in localization and second in classification in the ILSVRC 2014 ImageNet Challenge [39]. The VGG network is known for its simple architecture and competitive performance. The convolutional kernels each have a 3×3 pixel spatial extent. The nonlinear activation function applied after each convolution is a rectifier (*ReLU*), $f(x) = \max(0, x)$. The convolutional layers are applied in a series of blocks, and between the blocks, pooling layers are applied (in the original network, Max pooling, but here as in Refs. 37 and 38 we use Mean pooling). Blocks one and two contain two convolutions each, and blocks three, four, and five contain four convolutions each. The final stage of the network adds three fully connected layers—these do not directly encode spatial information, and so are not used for translation invariant characterization of images. We used the optimizing compiler Theano [50] and the neural network library Lasagne [51] to implement the CNN methods used in this paper.

III. METHODS

A. CNN texture vector representation of image statistics

Gatys *et al.* [37, 38] have developed a robust algorithm for statistical analysis of texture images relevant to materials microstructure, which we demonstrated in Fig. 1. Given an input texture image, the Gatys procedure extracts a texture vector from activations in the convolutional layers. This calculation is summarized in Fig. 2. Activations in the CNN are denoted by F_{ij}^l , where l is the layer index, i is the feature index, and j is the pixel index. At each layer, the Gram matrix captures correlations between feature i and feature k ,

$$G_{ik}^l = \sum_j F_{ij}^l F_{kj}^l. \quad (1)$$

The summation over pixel index j encodes invariance to translations in the image plane (up to boundary effects).




FIG. 2: A schematic view of the creation of a texture vector representation of an image using a CNN. The input image is processed upward through multiple CNN layers. Activations at each layer produce a stack of filtered images. Correlations between filtered images, averaged over the image plane, are collected. These statistics are then concatenated (with optional weighting) to form a texture vector. The activations at higher layers capture texture information over larger spatial scales.

Compared to the mean feature activations $\sum_j F_{ij}^l$ alone, the Gram matrix offers much richer statistical information [37]. In the following we suppress feature indices i and k and view the Gram matrix, $G_{ik}^l \rightarrow G^l$, as a summary of activation statistics on layer l . For the purposes of texture synthesis, Gatys et al. introduce a scalar, positive-definite loss between two images x_i and x_j (let us note explicitly that i and j now index images) with Gram matrices $G^l(x_i)$ and $G^l(x_j)$:

$$L^l(x_i, x_j) = \frac{1}{A_l} \|\hat{G}^l(x_i) - \hat{G}^l(x_j)\|^2, \quad (2)$$

with $\|\cdot\|$ the Euclidean norm. $A_l = 4N_l^2 M_l^2$ is a normalization factor for the loss on layer l with N_l features and M_l pixels. The total loss is the weighted sum of layer-wise losses,

$$L(x_i, x_j) = \sum_l w_l L^l(x_i, x_j). \quad (3)$$

In this work we use the VGG network [28], normalized as in [37, 38], and apply equal weight ($w_l = 1$) to each of the following layers: “conv1.1”, “conv2.1”, “conv3.1”, “conv4.1”, and “conv5.1”. It is convenient to define rescaled Gram matrices, $\hat{G}^l = (w_l/A_l)G^l$. Their concatenation $\hat{G} = (\hat{G}^1, \hat{G}^2, \dots)$ is the scaled texture vector. The layers we use have feature sizes of $[N_1, N_2, \dots] =$

$[64, 128, 256, 512, 512]$, resulting in a total texture vector length of $\sum_l N_l^2 \approx 5 \cdot 10^5$ elements. From the texture vector of two images, we may form a distance between images

$$d_{ij}^2 = \|\hat{G}(x_i) - \hat{G}(x_j)\|^2 = L(x_i, x_j). \quad (4)$$

That is, \hat{G} as a function endows two images x_i and x_j with a Euclidean distance d_{ij} based on their texture representations within the CNN. We will show that this distance is a useful input to machine learning algorithms, and in particular, manifold learning (see Sec. III C).

B. Power spectrum statistics

To benchmark the CNN texture vector representation, we compare it against the power spectrum (PS) associated with two-point correlations in the image. This approach is commonly employed for statistical characterization of microstructure. Our test dataset contains single-component (grayscale) images x_i , each represented as a scalar field $\phi_i(\mathbf{r})$. Assuming translation invariance, the two-point correlation function of $\phi(\mathbf{r})$ is

$$P_2(\Delta\mathbf{r}) = \int \phi(\mathbf{r}')\phi(\mathbf{r}' + \Delta\mathbf{r}) \, d\mathbf{r}'. \quad (5)$$

If ensemble averaged, the full set of n -point correlation functions would capture all information about the statistical distribution of images.

The PS (also known as the structure factor) $S(\mathbf{q})$ is the Fourier transform of P_2 , which can be expressed as

$$S(\mathbf{q}) = \tilde{\phi}(\mathbf{q})\tilde{\phi}(-\mathbf{q}) = |\tilde{\phi}(\mathbf{q})|^2, \quad (6)$$

where $\tilde{\phi}(\mathbf{q})$ is the Fourier transform of $\phi(\mathbf{r})$. For this analysis, we compute the PS after rescaling $\phi(\mathbf{r})$ to the range $[-1, 1]$.

To develop low-dimensional representations of microstructures, we require a distance between microstructures. Given two images x_i and x_j and their respective power spectra $S_i(\mathbf{q})$ and $S_j(\mathbf{q})$, we obtain a new distance d_{ij} between the images,

$$d_{ij}^2 = \int [S_i(\mathbf{q}) - S_j(\mathbf{q})]^2 d\mathbf{q}, \quad (7)$$

which should be contrasted with the CNN distance in Eq. (4).

C. Manifold Learning with Multidimensional Scaling

To assess the quality of the texture vectors for characterizing images we perform manifold learning. The goal of manifold learning is to find a low-dimensional representation of the data that faithfully captures distance information [here, from Eq. (4) or (7)]. Multidimensional Scaling [52–54] (MDS) implements this principle as follows: Given a dataset $\{x_i\}$, a distance function $d_{ij} = d(x_i, x_j)$ between datapoints, and an embedding dimension D^* , the goal is to find embeddings $x_i \mapsto \hat{x}_i \in \mathbb{R}^{D^*}$ such that the embedded Euclidean distance $\hat{d}_{ij} = \|\hat{x}_i - \hat{x}_j\|$ best matches d_{ij} , *i.e.* minimizes a stress function σ . Here, we use Kruskal’s stress [52],

$$\sigma = \sqrt{\frac{\sum (d_{ij} - \hat{d}_{ij})^2}{\sum \hat{d}_{ij}^2}} \quad (8)$$

We will consider embedding dimensions $D^* \lesssim 10$, which is very small compared to the dimension of the full CNN texture vector, roughly 2.5×10^5 (Sec. III A).

Note that MDS seeks \hat{d}_{ij} that globally matches distances d_{ij} , and thus captures information about *all* image pairs. Other schemes, such as Local Linear Embedding [55] or Isomap [56], instead work with a local sparsification of the distance matrix d_{ij} . In this work, we select MDS because of its direct interpretability and conceptual simplicity. MDS requires only one hyperparameter, the embedding dimension D^* . We use the scikit-learn [57] implementation of MDS, which applies an iterative majorization algorithm [58] to optimize the embedding stress σ , Eq. (8).




FIG. 3: The space of synthetic textures, generated by tunable parameters A , k , and θ .

IV. TASKS

A. Image generation process

We argue that although the space of materials microstructure is very rich, it will admit an effective low-dimensional representation. For example, a description of the materials processing (*e.g.* composition, thermodynamic variables and their rates of change) should be more compact than a direct description of the resulting microstructure. A statistical analysis of microstructure is valuable in that it may lead to further dimensionality reduction; multiple different processing paths may lead to the same microstructure. In this work, we study a database of synthetic 2-D microstructure images generated from a stochastic process with a few tunable generating variables. We use Perlin noise [59] to generate marble-like stochastic images. The method procedurally calculates a smooth multi-scale noise function $\mathbf{h}(\mathbf{r})$ that generates distorted spatial points $\mathbf{u}_A = \mathbf{r} + A\mathbf{h}(\mathbf{r})$ with noise amplitude A . Then each texture image x_i is realized as a 2-D scalar field

$$\phi_i(\mathbf{r}) = \cos[2\pi k_i \mathbf{u}_{A_i} \cdot \hat{\mathbf{n}}_{\theta_i}], \quad (9)$$

where $\hat{\mathbf{n}}_{\theta}$ is a unit vector with angle θ . That is, each image x_i consists of sinusoidal oscillation parameterized by angle (θ_i), scale (k_i), and noise amplitude (A_i) parameters. This three-dimensional parameter space is shown in Fig. 3.

Figure 4 illustrates the multi-scale nature of the set $\{x_i\}$ of stochastic texture images. At small noise values, the power spectrum is peaked on two Fourier modes. With increasing noise amplitudes, the peaks of the power spectrum broaden.




FIG. 4: Row (a): Synthetic microstructures with scale parameter $k = 15$ and varying noise amplitudes A . Row (b): Associated power spectrum, zoomed to relevant region of Fourier space. To aid visualization, the intensities are scaled by factors of approximately 1, 3, and 100, such that their maxima appear equally dark.

B. Angle reconstruction task

Our first task is to reconstruct a 1-D manifold of images of fixed noise amplitude and scale parameter, but varying angle. For each trial, the scale parameter was fixed to $k = 15$, corresponding to a modulation wavelength of $1/15$ in units of the linear system size. The angles θ_i take values $(i/N)\pi$ for $i \in \{0, 1, \dots, N-1\}$. Note that $0 \leq \theta < \pi$ without loss of generality because $\hat{\mathbf{n}}_\theta = -\hat{\mathbf{n}}_{\theta+\pi}$ and thus θ and $\theta + \pi$ are equivalent for our textures, Eq. (9). In this subsection we explore datasets with varying dataset sizes N and noise amplitudes A . We compute distances between the images via the CNN (Sec. III A) and PS (Sec. III B) methods, then use MDS (Sec. III C) to map the images into a $D^* = 2$ embedding space.

We quantify reconstruction quality as follows: First, we find the center of mass of all points in the embedded space, and use this as the origin. Second, we calculate angles φ_i about the origin, which are unique up to a single additive constant c . Finally, we seek a correspondence between the generating angles θ_i and the learned values $\varphi_i/2$. The factor of $1/2$ is necessary because θ_i ranges from 0 to π whereas φ_i ranges from 0 to 2π . We select

the constant c to minimize the root-mean-square error,

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_i [\theta_i - (\varphi_i + c)/2]^2} \quad (10)$$

Once c is optimized, we use the RMSE to measure the reconstruction quality.

Figure 5 shows embedded manifolds and corresponding angle reconstructions using dataset size $N = 50$ and noise amplitudes $A \in \{1, 2.5, 4\}$. For low noise amplitudes $A = 1$, the CNN distances produce a ring structure which reflects the generating angles (and associated periodicity) quite well, whereas the PS method fails. For intermediate $A = 2.5$, both CNN and PS distances generate good angle reconstructions, but there is much less scatter in the CNN embeddings. For large $A = 4$, the CNN continues to give good angle reconstructions despite scatter in the embedded points, whereas the PS method again fails. Note that, by construction, the PS method is rotationally symmetric, whereas the CNN method encodes rotational symmetry only approximately. Consequently, the CNN embeddings are somewhat elliptical.

Figure 6 shows the RMSE, Eq. (10), for the angle reconstruction task using a variety of dataset sizes N and noise amplitudes A . The CNN embeddings reliably reconstruct the generating angles θ_i for a wide range of N

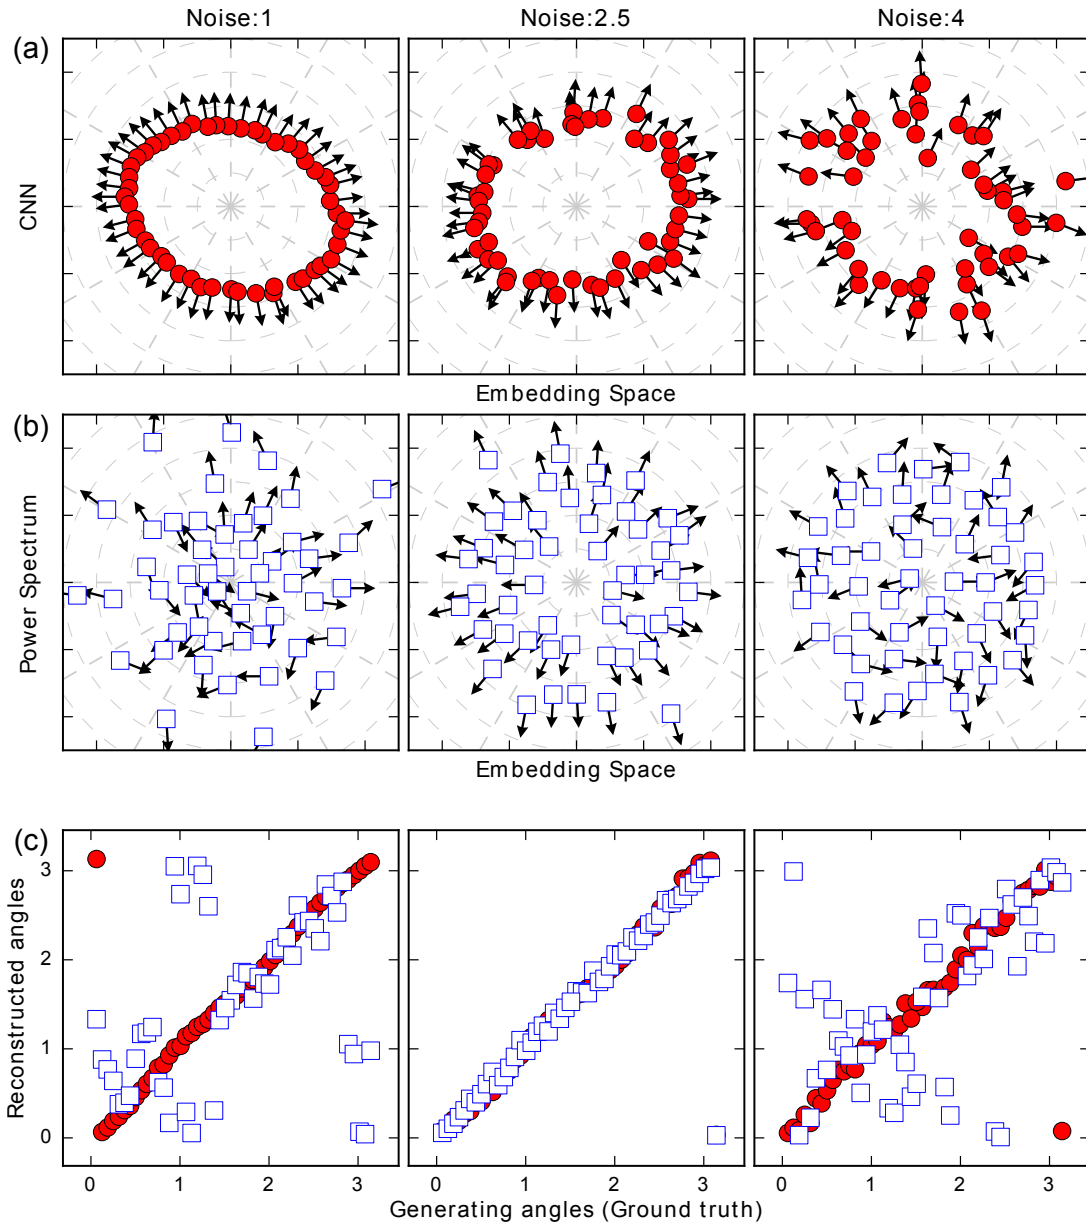


FIG. 5: (Color online) Angle reconstructions. Each dataset contains $N = 50$ images with varying angles θ_i , fixed scale $k = 15$, and fixed noise amplitude $A \in \{1, 2.5, 4\}$. Red circles correspond to the CNN method (Sec. III A) and blue squares correspond to the PS method (Sec. III B). Row (a): MDS embedded points \hat{x}_i using the CNN method. Arrows represent $2\theta_i$ where θ_i is the angle used to generate image angles. Row (b): MDS embedded points \hat{x}_i using the PS method. Row (c): Comparison of θ_i with corresponding angle $(\varphi_i + c)/2$ reconstructed from the embedding space. The CNN method yields excellent agreement, and clearly outperforms the PS method.

and A . However, the PS embeddings reconstruct θ_i only for a narrow window of A , and require a much larger N to reach comparable accuracy. This behavior can be understood by referring to Fig. 4: At very small A , the PS peaks are sharp, and there is little overlap between texture images with different angles. At very large A , the PS peaks broaden and exhibit great stochastic fluctuation. The best reconstructions occur at intermediate A , for which the peaks have some width but are not dominated by fluctuations, such that PS distances can accurately

capture differences in the angle parameter.

Figure 7 shows the embedding stress σ , Eq. (8), a measure of the fidelity of the MDS embedding. The CNN embedding exhibits low stress across a wide range of noise amplitudes, whereas the PS distances do not easily embed into a $D^* = 2$ embedding space. The stress of both CNN and PS embeddings grows with the noise amplitude. We interpret this as follows: At zero noise, the space of texture images has a single parameter, the angle. With finite noise, this 1-D manifold of texture im-




FIG. 6: (Color online) The angle reconstruction error, Eq. (10), as a function of noise amplitude A , with scale $k = 15$. Each dataset consists of N images with varying angle θ_i . The CNN method performs well across a wide range of noise amplitudes $A \lesssim 4.5$, whereas the PS method does best in a narrow range $2 \lesssim A \lesssim 3.5$. Each RMSE estimate represents an average over 100 independent trials.

ages expands into a much higher-dimensional space. The effective expansion volume increases monotonically with the noise amplitude. Consequently, it becomes increasingly difficult to embed this very high-dimensional manifold using $D^* = 2$, which is reflected in the increasing embedding stress.

C. Three dimensional manifold reconstruction task

Here we embed texture images from all three generating parameters shown in Fig. 3: the angle $0 \leq \theta < \pi$, scale $5 \leq k < 15$, and noise amplitude $0.5 \leq A < 2$. We generate a dataset of $N = 1000$ texture images by varying each parameter through 10 equally spaced increments. As before, we determine the distances between images using CNN (Sec. III A) and power spectrum (Sec. III B) methods, then use MDS (Sec. III C) to embed these distances into spaces of varying dimension D^* .

Figure 8 shows the embedding stress σ as a function of embedding dimension D^* . We observe a much lower stress using the CNN distances. The stress σ decays exponentially up to about $D^* = 6$ and flattens soon after. That is, with ≈ 6 descriptors per image, MDS has learned a representation of the texture images quite faithful to the CNN distances. Conversely, for the PS method, σ decays very slowly with D^* , suggesting that there is no natural low-dimensional embedding manifold.

The panels in Fig. 9 show the CNN method embeddings in a $D^* = 3$ space. The generating parameters A ,




FIG. 7: (Color online) The embedding stress σ as a function of noise amplitude A , with scale $k = 15$. Each dataset consists of N images with varying angle θ_i . The CNN method yields much lower stress embeddings than the PS method.




FIG. 8: (Color online) The stress σ as a function of embedding dimension D^* for a dataset of $N = 1000$ images with varying A , k , and θ . Stress from the CNN method (red) decreases approximately exponentially for $D^* \lesssim 6$. Stress from the PS method (blue) decreases much more slowly. That is, the PS distances do not naturally embed into low dimensions.

k , and θ emerge as a nonlinear coordinate system that spans a roughly conical solid in the embedding space. Column (a) of Figure 9 shows surfaces in which the noise amplitude A is held fixed, and the parameters θ and k are allowed to vary. For each A , the embedded points form an approximately conic surface. Cones with larger A are nested inside of cones with smaller A . Column (b) of Figure 9 shows surfaces in which the scale k is held

Embedding dimension D^*	2	3	4	5	6	10	50
R^2 for scale k	.532	.720	.901	.916	.916	.930	.980
R^2 for noise amplitude A	.183	.231	.908	.950	.951	.972	.983

TABLE I: The coefficients of determination R^2 for linear regression models that map embedded points \hat{x}_i to k_i or A_i . For embedding dimensions $D^* \geq 4$, the linear models achieve R^2 values near the ideal of 1.

fixed, allowing A and θ to vary. Surfaces of constant k are harder to describe, but resemble fragments of conic surfaces with varying angles. Surfaces with smaller k are nested inside of ones with larger k . Column (c) of Figure 9 shows surfaces in which the angle θ is held fixed, allowing A and k to vary. Rotating to a top-down view (Fig. 10), one observes that θ is very well captured by the azimuthal angle of the cone structure.

Figure 11 shows a 3-D projection of the $D^* = 4$ embedding. In this projection, we observe that the generating parameters appear approximately as cylindrical coordinates. The noise parameter A appears approximately linearly as a longitudinal coordinate.

To quantitatively evaluate the quality of embeddings for general dimension D^* , we performed linear regression to model the parameters A_i and k_i for each embedding point \hat{x}_i :

$$\hat{A}_i = \beta_A \cdot \hat{x}_i + \gamma_A, \quad (11)$$

$$\hat{k}_i = \beta_k \cdot \hat{x}_i + \gamma_k. \quad (12)$$

The regression vectors β_p and scalars γ_p for $p \in \{A, k\}$ were found using ordinary least squares by minimizing $\sum_i (p_i - \hat{p}_i)^2$. We then assess fit quality using the coefficient of determination,

$$R^2 = 1 - \frac{\sum_i (p_i - \hat{p}_i)^2}{\sum_i (p_i - \bar{p})^2}, \quad (13)$$

where $\bar{p} = \sum_i p_i / N$. An R^2 value of 1 indicates a perfectly linear relationship. Table I shows R^2 values for k and A in various embedding dimensions. In particular, $R^2 \approx 1$ is achieved already with $D^* = 4$ and higher-dimensional embeddings ($D^* > 4$) yield only marginal increase in fit quality.

Figure 12 shows the $D^* = 6$ embedding projected onto the two dimensions, $\{\beta_k, \beta_A\}$, that best linearly model the noise amplitude and scale parameters. The scale and noise parameters are nearly orthogonal; the angle between β_k and β_A is

$$\theta = \arccos\left(\frac{\beta_A \cdot \beta_k}{\|\beta_A\| \|\beta_k\|}\right) = 91.0^\circ. \quad (14)$$

V. DISCUSSION

The effectiveness of the PS on the angle reconstruction task (see Fig. 6) is best for moderate values of the

noise, but weak outside of this window. Although we discussed the mechanisms at play, this result is at first counterintuitive; The images consist of perturbed sinusoidal stripes that coincide well the Fourier basis used by the PS. The performance of the CNN texture vector is much better, achieving high accuracy and consistent performance across the spectrum of noise. This can be attributed to several advantages of the CNN-based approach.

First, the CNN uses local filters as opposed to global modes. Global features can suffer from interference effects, where similar small scale features which appear at a large distances from each other can add together destructively. Local features do not suffer from this type of failure mode, and so are more robust to noisy variations in patterns. This is similar to advantages of compact support in wavelet approaches to signal processing, which are well studied [60].

Second, compositions of convolutional filters and non-linear activations represent very non-trivial correlations between the pixels within their receptive fields, so that individual neurons are sensitive to higher order statistics that are not captured in the PS representation. For example, higher order statistics can directly characterize complex features such as domain edge curvature.

Lastly, the pooling layers in the CNN operate similarly to coarse-graining in physics, and is designed to capture relevant system characteristics while discarding unimportant ones [61]. In the CNN, repeated convolutional and pooling operations effectively implement coarse-graining over multiple layers of abstraction. Features that appear in deeper layers (*i.e.* further from the input) of the CNN have a spatially larger receptive field, and are more robust to small changes of the input due to the coarse-graining. Thus, larger-scale CNN features are naturally insensitive to smaller-scale texture details, which we believe is key to microstructure analysis as well as computer vision tasks. A trade-off with deep neural networks is that it can be difficult to understand concretely what a particular activation in a CNN represents; however, this is an active area of research [62–64].

VI. CONCLUSIONS AND FUTURE DIRECTIONS

We have introduced a method for unsupervised detection of the low-dimensional structure of a distribution of texture images using CNNs. We discuss the uses of this as a framework for the analysis of materials microstructure to learn dimensionality and topology of microstructure families using low-dimensional quantitative descriptions of microstructure. Compact microstructure characterization forms a platform for the construction of reduced order models that connect processing to microstructure, and microstructure to properties. This approach is applicable to small data sets, which is an important design factor in materials science and other disciplines where

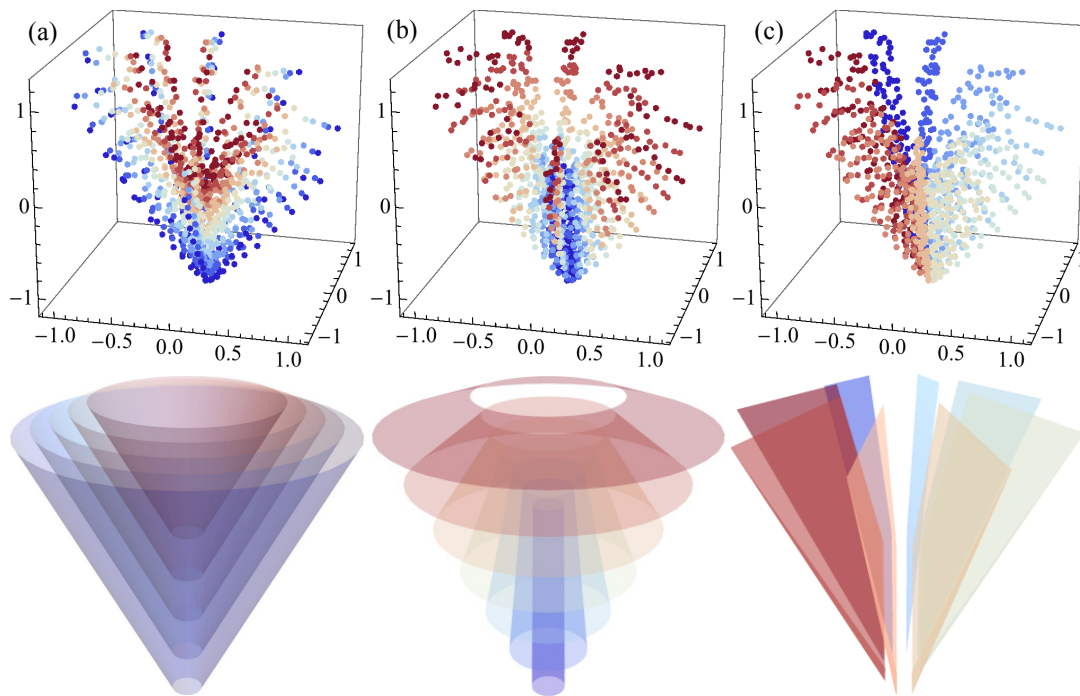


FIG. 9: (Color online) The $D^* = 3$ embedding of a dataset with $N = 1000$ image of varying noise amplitude A , scale k , and angle θ . Columns (a), (b), and (c) are colored by A , k , and θ , respectively. For each column, the top panel shows the embedded points, and the bottom panel displays a schematic representation of surfaces with constant generating parameter.




FIG. 10: (Color online) Axis-aligned view of the $D^* = 3$ embedding colored by angle θ (cf. Fig. 9), demonstrating strong correspondence with angles in the embedding space.

acquiring data can be expensive. In this work, we apply manifold learning to a synthetic dataset. This controlled context enables us to quantify the success of manifold learning. We anticipate that similar manifold learn-

ing approaches will prove effective for follow-on studies of real materials. For example, DeCost *et al.* recently demonstrated success in mapping the microstructures of ultra high carbon steels [40].

The method presented in this work is computationally efficient. In our Theano implementation, running on a single GPU, it takes about a millisecond to compute the distance between texture vectors that represent two images. MDS operates on all distance pairs, and thus scales quadratically with the size of the dataset. The MDS calculation on our full dataset of 10^3 synthetic images (Sec. IV C) completed in about 30 minutes. The dominant cost was the MDS embedding procedure, which took about 18 minutes. Calculating the $\approx 10^6/2$ texture vector distances took about 12 minutes.

A limitation of the transfer learning approach is that it requires a well-trained CNN with applicability to the target domain, which presently limits our analysis to 2-D micrographs. One path for improvement is to directly train CNNs on a large database of standardized microstructure images. Such a database could also be used to develop latent variable models (e.g. [65]) that would reflect the microstructural generation process. These end-to-end models would enable direct inference of low-dimensional generating parameters and direct generation of new microstructure image samples.

The work of Ref. 66 suggests that, instead of using transfer learning on natural images, it may be possible to characterize microstructure textures using randomized




FIG. 11: (Color online) The $D^* = 4$ embedding of a dataset with $N = 1000$ images of varying noise amplitude A , scale k , and angle θ . Panels (a), (b), and (c) are colored by A , k , and θ , respectively. To visualize the data, we select a 3-D projection that illustrates a decoupling of the generating parameters, in which they manifest as roughly cylindrical coordinates: A maps to longitudinal height, k to radius, and θ to azimuthal angle.

CNNs. Specifically, Ustyuzhaninov *et al.* conclude that suitably structured random, shallow, multiscale networks can, in some cases, be used to generate higher quality textures than those generated from a trained CNN. However, Fig. 1 of Ref. 66 shows that the distance matrix generated from the trained CNN is closer to the identity compared to the distance matrix generated from the random network. This suggests that the trained CNN is a better starting point for comparing texture images. The capability of random networks to perform low-dimensional embeddings of microstructures remains an open question. The use of random networks suggests exciting opportunities to operate on other data modalities, *e.g.* three-dimensional microstructure data [25, 41, 67] and/or grain orientation data [47, 68, 69], both of which are outside the domain of natural image characterization.

Lastly, we consider the rotation group, which factors into microstructure analysis in at least two ways. Firstly, rotations appear through spatial transformations of the image plane. The standard CNN architecture does not explicitly incorporate such transformations. This is evident in the small but consistent biases in the angular reconstruction task (see Figs. 5 and 6). However, the relatively strong performance of the network on this task indicates that the network has implicitly learned approximate representations of the rotation group during the training procedure. A second way that the rotation group appears is in grain orientation data. CNNs designed to process RGB images do not directly represent the group structure of crystalline orientations. Further texture characterization work might explicitly incorporate the action of rotations, *e.g.* building upon Refs. 70–74.

ACKNOWLEDGMENTS

We acknowledge funding support from a Laboratory Directed Research and Development (LDRD) DR (#20140013DR), and the Center for Nonlinear Studies (CNLS) at the Los Alamos National Laboratory (LANL). We also thank Prasanna Balachandran and James Theiler for useful discussions and feedback.

Appendix A: Attribution of microstructure images

The images appearing in Fig. 1 are released in the public domain, and available online. Permanent links are as follows. Top row, from left to right:

1. <https://commons.wikimedia.org/w/index.php?title=File:Gailbach-Tonalit.jpg&oldid=179586143>.
2. <https://commons.wikimedia.org/w/index.php?title=File:PearliteSph3.jpg&oldid=149930090>.
3. <https://commons.wikimedia.org/w/index.php?title=File:Martensit.jpg&oldid=144193548>.
4. <https://commons.wikimedia.org/w/index.php?title=File:Lamine316L.jpg&oldid=66965939>.

Bottom row, from left to right:

5. https://commons.wikimedia.org/w/index.php?title=File:Microstructure_of_rolled_and_annealed_brass;_magnification_400X.jpg&oldid=144992203.
6. <https://commons.wikimedia.org/w/index.php?title=File:Ferrite-perlite-steel-A285.jpeg&oldid=140529933>.




FIG. 12: (Color online) The 2-D projection, colored by (a) scale and (b) noise amplitude, of the $D^* = 6$ embedding of a dataset with $N = 1000$ images of varying A , k , and θ . This projection was selected using linear regression to find the two dimensions that best capture the scale k and noise amplitude A parameters.

7. <https://commons.wikimedia.org/w/index.php?title=File:Pearlite1.jpg&oldid=184321025>.

8. https://commons.wikimedia.org/w/index.php?title=File:Feuerverzinkte_Oberfl%C3%A4che.jpg&oldid=140580291

-
- [1] H. Kumar, C. Briant, and W. Curtin, *Mech. Mater.* **38**, 818 (2006).
- [2] M. Ostoja-Starzewski, *Microstructural randomness and scaling in mechanics of materials* (CRC Press, 2007).
- [3] M. Wang and N. Pan, *Mater. Sci. Eng. R-Rep.* **63**, 1 (2008).
- [4] D. T. Fullwood, S. R. Niezgod, B. L. Adams, and S. R. Kalidindi, *Prog. Mater. Sci.* **55**, 477 (2010).
- [5] S. Torquato, *Annu. Rev. Mater. Res.* **40**, 101 (2010).
- [6] K. Rajan, *Informatics for materials science and engineering: data-driven discovery for accelerated experimentation and application* (Butterworth-Heinemann, 2013).
- [7] S. R. Kalidindi, *Int. Mater. Rev.* **60**, 150 (2015).
- [8] T. Lookman, F. J. Alexander, and K. R. (Eds.), *Information Science for Materials Discovery and Design*, edited by K. R. Turab Lookman, Francis J. Alexander (Springer International Publishing, 2016).
- [9] P. M. Voyles, *Curr. Opin. Solid State Mater. Sci.* (2016), in press.
- [10] D. Xue, P. V. Balachandran, J. Hogden, J. Theiler, D. Xue, and T. Lookman, *Nat. Commun.* **7** (2016).
- [11] S. R. Kalidindi, S. R. Niezgod, and A. A. Salem, *JOM* **63**, 34 (2011).
- [12] Y. Liu, M. S. Greene, W. Chen, D. A. Dikin, and W. K. Liu, *Comput. Aided Des.* **45**, 65 (2013).
- [13] S. R. Niezgod, A. K. Kanjarla, and S. R. Kalidindi, *Integr. Mater. Manuf. Innov.* **2**, 3 (2013).
- [14] Y. Jiao, F. H. Stillinger, and S. Torquato, *Phys. Rev. E* **76**, 031110 (2007).
- [15] D. T. Fullwood, S. R. Niezgod, and S. R. Kalidindi,

- Acta Mater. **56**, 942 (2008).
- [16] Y. Jiao, F. H. Stillinger, and S. Torquato, Proc. Natl. Acad. Sci. USA **106**, 17634 (2009).
- [17] D. D. Chen, Q. Teng, X. He, Z. Xu, and Z. Li, Phys. Rev. E **89**, 013305 (2014).
- [18] Y. Jiao, F. H. Stillinger, and S. Torquato, Phys. Rev. E **81**, 011105 (2010).
- [19] Y. Jiao, F. H. Stillinger, and S. Torquato, Phys. Rev. E **82**, 011106 (2010).
- [20] C. J. Gommès, Y. Jiao, and S. Torquato, Phys. Rev. E **85**, 051140 (2012).
- [21] C. J. Gommès, Y. Jiao, and S. Torquato, Phys. Rev. Lett. **108**, 080601 (2012).
- [22] S. Niezgodá, D. Fullwood, and S. Kalidindi, Acta Mater. **56**, 5285 (2008).
- [23] C. E. Zachary and S. Torquato, Phys. Rev. E **84**, 056102 (2011).
- [24] K. M. Gerke, M. V. Karsanina, R. V. Vasilyev, and D. Mallants, Europhys. Lett. **106**, 66002 (2014).
- [25] H. Xu, D. A. Dikin, C. Burkhardt, and W. Chen, Comput. Mater. Sci. **85**, 206 (2014).
- [26] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Proc. IEEE **86**, 2278 (1998).
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Adv. Neural Inf. Process. Syst., 1097 (2012).
- [28] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” (2014), arXiv:1409.1556 [cs.CV].
- [29] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” (2014), arXiv:1409.4842 [cs.CV].
- [30] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” (2015), arXiv:1512.03385 [cs.CV].
- [31] Y. LeCun, Y. Bengio, and G. E. Hinton, Nature **521**, 436 (2015).
- [32] T. Hao, T. Raiko, A. Ilin, and J. Karhunen, in *Artif. Neural Networks Mach. Learn. – ICANN 2012, Lecture Notes in Computer Science*, Vol. 7553, edited by A. E. Villa, W. Duch, P. Érdi, F. Masulli, and G. Palm (Springer Berlin Heidelberg, 2012) pp. 124–131.
- [33] J. J. Kivinen and C. Williams, in *International Conference on Artificial Intelligence and Statistics* (2012) pp. 638–646.
- [34] H. Luo, P. L. Carrier, A. Courville, and Y. Bengio, Proc. 16th Int. Conf. Artif. Intell. Stat. **31**, 415 (2013).
- [35] Q. Gao and S. Roth, in *Struct. Syntactic, Stat. Pattern Recognit.*, Lecture Notes in Computer Science, Vol. 8621, edited by P. Fránti, G. Brown, M. Loog, F. Escolano, and M. Pelillo (Springer Berlin Heidelberg, 2014) pp. 434–443.
- [36] L. Theis and M. Bethge, in *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., 2015) pp. 1927–1935.
- [37] L. A. Gatys, A. S. Ecker, and M. Bethge, “Texture synthesis using convolutional neural networks,” (2015), arXiv:1505.07376 [cs.CV].
- [38] L. A. Gatys, A. S. Ecker, and M. Bethge, “A neural algorithm of artistic style,” (2015), arXiv:1508.06576 [cs.CV].
- [39] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, Int. J. Comput. Vision **115**, 211 (2015).
- [40] B. L. DeCost, T. Francis, and E. A. Holm, Acta Mater. **133**, 30 (2017).
- [41] V. Sundararaghavan and N. Zabaras, Comput. Mater. Sci. **32**, 223 (2005).
- [42] B. L. DeCost and E. A. Holm, Comput. Mater. Sci. **110**, 126 (2015).
- [43] S. V. Kalinin, B. G. Sumpter, and R. K. Archibald, Nat. Mater. **14**, 973 (2015).
- [44] R. Liu, A. Kumar, Z. Chen, A. Agrawal, V. Sundararaghavan, and A. Choudhary, Sci. Rep. **5**, 11551 (2015).
- [45] R. Bostanabad, A. T. Bui, W. Xie, D. W. Apley, and W. Chen, Acta Mater. **103**, 89 (2016).
- [46] A. Chowdhury, E. Kautz, B. Yener, and D. Lewis, Comput. Mater. Sci. **123**, 176 (2016).
- [47] A. D. Orme, I. Chelladurai, T. M. Rampton, D. T. Fullwood, A. Khosravani, M. P. Miles, and R. K. Mishra, Comput. Mater. Sci. **124**, 353 (2016).
- [48] K. He, X. Zhang, S. Ren, and J. Sun, in *The IEEE International Conference on Computer Vision (ICCV)* (2015) pp. 1026–1034.
- [49] I. Goodfellow, Y. Bengio, and A. Courville, “Deep learning,” (2016), Book in preparation for MIT Press.
- [50] Theano Development Team, “Theano: A Python framework for fast computation of mathematical expressions,” (2016), arXiv:1605.02688 [cs.SC].
- [51] S. Dieleman, J. Schluiter, C. Raffel, E. Olson, S. K. Sønderby, D. Nouri, D. Maturana, M. Thoma, E. Battenberg, J. Kelly, *et al.*, “Lasagne: First release.” (2015).
- [52] J. B. Kruskal, Psychometrika **29**, 1 (1964).
- [53] I. Borg and P. J. F. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer Series in Statistics (Springer-Verlag New York, 2005).
- [54] J. Franklin, Math. Intell. **27**, 83 (2008).
- [55] S. T. Roweis and L. K. Saul, Science **290**, 2323 (2000).
- [56] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, Science **290**, 2319 (2000).
- [57] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, J. Mach. Learn. Res. **12**, 2825 (2011).
- [58] J. De Leeuw, in *Recent Developments in Statistics*, edited by J. R. Barra, F. Brodeau, G. Romier, and B. van Cutsem (Verlag: North-Holland, 1977) pp. 133–145.
- [59] K. Perlin, SIGGRAPH Comput. Graph. **19**, 287 (1985).
- [60] S. Mallat, *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*, 3rd ed. (Academic Press, 2008).
- [61] An explicit correspondence between the variational renormalization group and deep learning using Boltzmann Machines is demonstrated in P. Mehta and D. J. Schwab, “An exact mapping between the Variational Renormalization Group and Deep Learning,” (2014), arXiv:1508.06576 [cs.CV].
- [62] D. Erhan, A. Courville, and Y. Bengio, *Understanding Representations Learned in Deep Architectures*, Tech. Rep. 1355 (Université de Montréal/DIRO, 2010).
- [63] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, “Understanding neural networks through deep visualization,” (2015), arXiv:1506.06579 [cs.CV].
- [64] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell,

- D. Parikh, and D. Batra, “Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization,” (2016), arXiv:1610.02391 [cs.CV].
- [65] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” (2013), arXiv:1312.6114 [stat.ML].
- [66] I. Ustyuzhaninov, W. Brendel, L. A. Gatys, and M. Bethge, “Texture synthesis using shallow convolutional networks with random filters,” (2016), arXiv:1606.00021 [cs.CV].
- [67] A. Rollett, S.-B. Lee, R. Campman, and G. Rohrer, *Annu. Rev. Mater. Res.* **37**, 627 (2007).
- [68] B. Adams, S. Wright, and K. Kunze, *Metall. Trans. A* **24**, 819 (1993).
- [69] F. J. Humphreys, *J. Mater. Sci.* **36**, 3833 (2001).
- [70] S. Mallat, *Commun. Pure Appl. Math.* **65**, 1331 (2012).
- [71] T. S. Cohen and M. Welling, “Transformation properties of learned visual representations,” (2014), arXiv:1412.7659 [cs.LG].
- [72] R. Gens and P. M. Domingos, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014) pp. 2537–2545.
- [73] S. Dieleman, J. D. Fauw, and K. Kavukcuoglu, “Exploiting cyclic symmetry in convolutional neural networks,” (2016), arXiv:1602.02660 [cs.LG].
- [74] T. S. Cohen and M. Welling, “Group equivariant convolutional networks,” (2016), arXiv:1602.07576 [cs.LG].